

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): Kamakshi Sridhar **Docket:** 135909
Serial No.: 10/020,593 **Art Unit:** 2616
Filed: 12/13/2001 **Examiner:** Derrick W. Ferris
Title: Load Balancing Technique for a Resilient Packet Ring

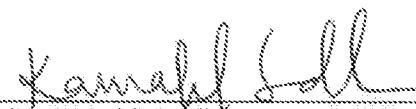
DECLARATION UNDER 37 CFR 1.131

I, Kamakshi Sridhar, do declare and say:

1. I am the inventor of the invention described in this patent application, residing at 2712 Oates Drive, Plano, TX 75093.
2. I conceived the invention prior to November 28, 2001, the filing date of Provisional application No. 60/333,936, and the filing date relied upon by U.S. Patent Application Pub. No. 2003/0103449 A1 of Barsheshet et al., cited as prior art in my application. My conception of the invention is evidenced by the description of the invention in the attached draft of my present patent application prepared by patent counsel and sent to me on November 23, 2001, along with the attached cover letter of that date.
3. After I received the attached application draft, I undertook a review of it, and patent counsel incorporated changes into the application.
4. The present application Serial No. 10/020,593 was filed Dec. 13, 2001.

I further declare that all statements made herein of my own knowledge are true, and that all statements made on information and belief are believed to be true; and further that all these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Date: June 5, 2006


Kamakshi Sridhar

SMITH, DANAMRAJ & YOUST

A PROFESSIONAL CORPORATION

ATTORNEYS AND COUNSELORS
INTELLECTUAL PROPERTY LAW AND RELATED MATTERS

STEVEN W. SMITH
SHREEN K. DANAMRAJ
LAWRENCE R. YOUST

12900 PRESTON ROAD, SUITE 1200, LB 15
DALLAS, TEXAS 75230-1328

TELEPHONE (972) 720-1202
FAX (972) 720-1139
EMAIL: SHREEN@SDYLAW.COM

November 23, 2001

Via Federal Express / Courier

CONFIDENTIAL AND PRIVILEGED

Ms. Kamakshi Sridhar
2712 Oates Drive
Plano, TX 75093

Re:	Patent Application For:	<i>LOAD BALANCING TECHNIQUE FOR A RESILIENT PACKET RING</i>
	Inventor(s):	Kamakshi Sridhar
	Our Ref. No.:	1285-0079US
	Alcatel Ref. No.:	135909

Dear Kamakshi:

Enclosed please find a draft of the above-referenced patent application prepared by my colleague, Brandi Sarfatis. Please go over the enclosed draft application, including the claims and drawings, to ensure that it represents a complete and technically accurate description of your invention.

Please be advised that if any minor errors of a typographical nature are discovered, they may be corrected in ink, provided you initial and date each such correction in the nearest marginal space of the pages where they are found.

If only minor errors are discovered, please correct them as set forth above, sign the enclosed ***Combined Declaration and Power of Attorney*** and ***Assignment*** forms, and return the entire package to me for filing with the U.S. Patent and Trademark Office.

If other serious errors or omissions are discovered, however, please let us know as soon as possible so that a revised draft with the necessary corrections may be prepared for your approval. The enclosed formal papers may be executed at that time.

Ms. Kamakshi Sridhar

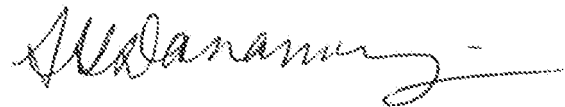
Re: 1285-0079US Draft Patent Application

Page TWO

Further, pursuant to Patent and Trademark Office rules, a statement must be submitted disclosing to the Office all presently known information believed material to the examination of the above-referenced patent application. Information is considered material if there is a substantial likelihood that the information would be considered important in deciding patentability. In this regard, please note that we have already received a packet of materials, presumably provided by you, that relate to the background of the subject matter covered in the present patent application. Nonetheless, if you happen to have in your possession additional materials such as patents, publications, or other published materials relating to the subject invention, I would appreciate receiving such materials at your earliest convenience. You are not required, however, to do any research to locate information that is not already known to you.

Accurate and complete disclosure of material information facilitates the examination phase of the application, and most importantly, strengthens the presumption of validity accorded any patent issuing thereon. If you have any questions concerning the submission of material information or the present patent application, please do not hesitate to call Brandi or me.

Yours very truly,



Shreen K. Danamraj

Enclosures

cc: Ms. Jessica Smith, Esq.
Intellectual Property Department
Alcatel USA

File copy

DRAFT

CONFIDENTIAL

PATENT APPLICATION
DOCKET NO.: 1285-0079US
ALC-135909

NOV 23 2001

"EXPRESS MAIL" Mailing Label No.....
Date of Deposit.....

LOAD BALANCING TECHNIQUE FOR A RESILIENT PACKET RING

CROSS-REFERENCE TO RELATED APPLICATION(S)

5 This application discloses subject matter related to
the subject matter disclosed in commonly owned, co-
pending U.S. Patent Application No.: _____
(Atty. Docket No. 1285-0078US), entitled "SIGNALING FOR
CONGESTION CONTROL, LOAD BALANCING, AND FAIRNESS IN A
RESILIENT PACKET RING", filed _____ in the
10 name(s) of: Kamakshi Sridhar.

BACKGROUND OF THE INVENTION

Technical Field of the Invention

15 The present invention generally relates to resilient
packet rings ("RPRs"). More particularly, and not by way
of any limitation, the present invention is directed to
a load balancing technique for use in an RPR.

Description of Related Art

20 Resilient Packet Ring ("RPR") is a packet-based
technology designed primarily for implementing
metropolitan area networks ("MANs"). The basic topology
of RPR is a dual counter-rotating ring in which both
rings transport both data and control packets. RPR

technology includes a Medium Access Control ("MAC") mechanism that controls access to the RPR.

FIG. 1 is an example RPR network 100 consisting of five nodes, respectively designated A-E, and two rings, illustrated in FIG. 1 as an outer ring X and an inner ring Y. Each node A-E is connected to its adjacent nodes via at least two links; one on the inner ring and one on the outer ring. Links on the inner ring are designated a-e, while links on the outer ring are designated α - ϵ . Accordingly, nodes A and B are connected by links α and e, nodes B and C are connected by links β and d, nodes C and D are connected by links γ and c, nodes D and E are connected by links δ and b, and nodes E and A are connected by links ϵ and a.

Information is carried in one and only one direction on a ring, either clockwise or counter-clockwise. As illustrated in FIG. 1, the outer ring X carries data and control packets in the clockwise direction and the inner ring Y carries data and control packets in the counter-clockwise direction. Features of RPR include spatial reuse and destination stripping, fast protection, media independent MAC, support for layer 2 devices (e.g., Ethernet switches) and layer 3 devices (e.g., routers).

RPR operates as a distributed switch at layer 2. Alternatives to RPR include SONET/SDH rings and Ethernet switches. The former, implemented with routers, is expensive and inefficient for dynamic traffic patterns; however, it provides efficient protection. The latter

offers no protection bandwidth or fairness, but does provide efficient bandwidth utilization. Moreover, in contrast to layer 3 devices, RPR has ring awareness. RPR, therefore, is intended to embody the best features of both SONET and Ethernet switches, that is, bandwidth efficiency, fairness, and fast protection. To realize these features, the RPR should provide a basic ring access mechanism, additional mechanisms to provide satisfactory performance for traffic transiting the ring, and appropriate protection features.

Several RPR terms and concepts will now be defined. Referring again to FIG. 1, assume that traffic destined for node D enters the ring at node A and passes through nodes B and C. Node A is referred to as the "ingress node", node D is referred to as the "egress node", and nodes B and C are referred to as "intermediate nodes." As previously noted, the fiber connecting two nodes is referred to as a "link". From the perspective of nodes B and C, traffic from node A is "ring traffic". Traffic entering node B or C is referred to as "tributary traffic," with respect to node A traffic.

At the ingress node, as soon as a flow is admitted to the RPR, an RPR header is created and attached to every packet within the flow. The RPR header contains a Class of Service ("CoS") marking, a preemption bit, a ring identifier, and a destination MAC address. Each intermediate node will examine the incoming packet and perform an address lookup. If the packet is destined for

the current node, it is taken off the ring at that node. If it is destined for some other node, it is placed back on to the ring.

5 One problem inherent in RPR architecture is the problem of load imbalance between two or more rings or wavelengths. Hereinafter, unless otherwise explicitly specified, the term "rings" will be deemed to encompass both rings and wavelengths. Load imbalance results from excessive traffic on one ring, as compared to little
10 traffic on another ring. In an RPR with two rings, there are only two options for sending traffic from one node to another node, either on one ring or on the other. In such cases, traffic is typically sent on the ring with the fewest hops. This may result in a situation where
15 a lot of traffic travels from one node to another on only one ring, because the ring has fewer hops as compared to the other ring.

Additionally, because RPR employs destination stripping, certain segments on one ring are typically
20 more heavily loaded as compared to other segments on the same ring. Also, certain segments on one ring may be much more heavily loaded as compared to segments on the other ring.

As previously indicated, load imbalance may be
25 defined as the concurrent overutilization of capacity on one or more links in the RPR and the underutilization of capacity of other links in the RPR. Load imbalance can be illustrated as follows. Referring again to FIG. 1,

traffic going from node A to node B can travel either on the outer ring, via link α , or on the inner ring via links a, b, c, and d. For obvious reasons, traffic will tend to take the link α , since it is the more direct route. In the absence of any other traffic on the RPR, the link α will be overutilized, while the other links will be underutilized. The solution is to "load balance" the traffic between the links so that packets of a given class will experience approximately the same delay regardless of what the shortest path is.

Thus, traffic going from one node to another experiences varying amounts of delay depending on how many and which heavily loaded segments it traverses. Since traffic performance is judged by total end-to-end delay, the traffic performance (delay, dropped packets, etc.) on one ring may be significantly worse than the traffic performance on the other ring for a single class of service. This results in congestion on one ring versus unused capacity on the other ring and therefore requires load balancing.

Load balancing in RPR involves two functions: (a) diverting incoming flows from more heavily loaded rings to alternate paths on other rings, where they will likely experience more hops, but better performance, due to less traffic; and (b) improving the performance of existing traffic on the heavily loaded ring. Perfect load balancing in RPR, for all classes, between all pairs of nodes, is challenging due to the ring structure and

destination stripping features of RPR. Load balancing performed for one pair of nodes for one class may result in traffic redistribution that might necessitate load balancing between another pair of nodes for the same or
5 different traffic class. Dynamic traffic patterns and time-dependent flow arrivals and departures also complicate load balancing in RPR.

Prior art for load balancing traffic in multiple rings (or wavelengths) of RPR is currently not known to
10 exist. Prior art for changing quality of service ("QoS") parameters of various classes relative to one another is inadequate to help divert traffic to alternate paths on multiple rings. Prior art to change the token bucket ("TB") parameters, including peak data rate and maximum
15 burst size, for aggregate traffic flows does not allow the behavior of traffic classes to be changed relative to one another. While it may be adequate to reduce congestion in a network, it is inadequate to implement load balancing between rings in an RPR. Finally, all
20 known prior art is restricted to only TB parameters and does not consider the interaction between the bandwidth broker ("BB") parameters, which include allocated bandwidth, available bandwidth, and used bandwidth, and the QoS parameters, which include the TB parameters as
25 well as the class based queuing ("CBQ") parameters.

SUMMARY OF THE INVENTION

Accordingly, the present invention advantageously provides a technique for load balancing in an RPR. In one embodiment, the present invention comprises
5 implementing on every node an QoS/BB monitor to achieve load balancing in an RPR. The present invention is applicable to dual-ring and multi-ring RPR, as well as wavelength division multiplexing ("WDM") and dense wavelength division multiplexing ("DWDM") RPR. The
10 QoS/BB monitor on a node is common to all rings in the RPR. It has knowledge of traffic performance for each class on each ring of the RPR, which information is obtained through periodic measurements or in response to failure events. This allows the monitor to vary the QoS
15 parameters on a node, for a particular traffic class, to achieve load balancing. Likewise, the QoS/BB monitor can signal to the BB entity at higher layers to vary the BB parameters on a node for a particular class to achieve load balancing.

20 Changing BB parameters affects all admission of new traffic flows. Changing QoS parameters affects performance of traffic already on the ring. BB parameter variations result in coarse variation of traffic performance, while QoS parameter variations are likely to
25 result in finer variation of traffic performance. Through the combination of varying QoS parameters and BB parameters, the QoS/BB monitor can throttle traffic on one ring and permit more traffic on another ring, for a

particular traffic class, to achieve load balancing. By signaling to the BB entity to change its parameters, the ability to divert traffic at the node ingress from one ring to another ring may be achieved.

5 The key benefit of having one entity that controls both the QoS parameters and the BB parameters is that the traffic behavior can be modified appropriately, in terms of direction and servicing rate, to ultimately achieve load balancing.

10 In one aspect, the invention comprises a method of implementing load balancing in an RPR network comprising a plurality of nodes and first and second rings each comprising a plurality of links for carrying information between the nodes in a clockwise direction and a
15 counterclockwise direction, respectively, wherein adjacent ones of the nodes are connected by two of the links. The method comprises the steps of: for one of the nodes, determining whether a load imbalance exists at the node a load imbalance in connection with a first class of
20 service; and responsive to a determination that a load imbalance exists, changing Bandwidth Broker ("BB") parameters at the node for the first class of service to cause new flows to be diverted from a more heavily loaded one of the rings to a less heavily loaded one of the
25 rings and changing Quality of Service ("QoS") parameters at the node for the first class of service to improve traffic performance on the more heavily loaded one of the

rings, while maximizing traffic performance on the less heavily loaded one of the rings.

In another aspect, the present invention comprises an apparatus for implementing load balancing in an RPR network comprising a plurality of nodes and first and second rings each comprising a plurality of links for carrying information between the nodes in a clockwise direction and a counterclockwise direction, respectively, wherein adjacent ones of the nodes are connected by two of the links. The apparatus comprises, at one of the nodes, detector means and/or its equivalents for detecting at the node a load imbalance in connection with a first class of service; a structure responsive to detection at the node of a load imbalance for changing BB parameters at the node for the first class of service to cause new flows to be diverted from a more heavily loaded one of the rings to a less heavily loaded one of the rings; and a structure responsive to detection at the node of a load imbalance for changing QoS parameters at the node for the first class of service to improve traffic performance on the more heavily loaded one of the rings, while maximizing traffic performance on the less heavily loaded one of the rings.

In a further aspect, the invention comprises an apparatus for implementing load balancing in an RPR network comprising a plurality of nodes and first and second rings each comprising a plurality of links for carrying information between the nodes in a clockwise

direction and a counterclockwise direction, respectively,
wherein adjacent ones of the nodes are connected by two
of the links. The apparatus comprises, at one of the
nodes, a QoS/BB monitor responsive to detection at the
5 node of a load imbalance in connection with a first class
of service for signaling to a BB to change BB parameters
at the node for the first class of service to cause new
flows to be diverted from a more heavily loaded one of
the rings to a less heavily loaded one of the rings and
10 for changing QoS parameters at the node for the first
class of service to improve traffic performance on the
more heavily loaded one of the rings, while maximizing
traffic performance on the less heavily loaded one of the
rings.

15
BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present
invention may be had by reference to the following
Detailed Description when taken in conjunction with the
20 accompanying drawings wherein:

FIG. 1 depicts an exemplary RPR network arrangement
in which teachings of the present invention may be
advantageously practiced;

FIG. 2 (Prior Art) depicts a block diagram of an
25 exemplary embodiment of a single node of an RPR network
according to teachings of the prior art;

FIG. 3 depicts a block diagram of an exemplary and
presently preferred embodiment of a single node of an RPR

network in which teachings of the present invention may be advantageously practiced; and

FIG. 4 depicts a flowchart of a load balancing technique for implementation in an RPR network in accordance with the teachings of the present invention.

DETAILED DESCRIPTION OF THE DRAWINGS

In the drawings, like or similar elements are designated with identical reference numerals throughout the several views thereof, and the various elements depicted are not necessarily drawn to scale. Referring again to FIG. 1, depicted therein is an exemplary RPR network 100 in which the teachings of the present invention may be advantageously practiced. It should be recognized that, although the RPR network 100 comprises five nodes A-E, in general, there can be more or fewer nodes on the network 100. Moreover, although each of the nodes A-E are illustrated as being connected to its adjacent nodes by only two links, it should be recognized that the principles of the invention described herein can be applied to an RPR comprising more than two rings, and therefore more than two links between adjacent nodes. The teachings of the present invention are also applicable to WDM and DWDM RPR networks.

Consider a large volume of traffic going from node A to node C. Two choices exist: (a) the path including links α and β (path α - β) via the outer (clockwise) ring; and (b) the path including links a, b, and c (path a-b-c)

via the inner (counter-clockwise) ring. Typical path selection protocols will result in the selection of the first path (path α - β), as it has the fewest number of hops. New traffic entering node A destined for node C
5 must be diverted to path a-b-c for better performance. This is achieved through the load balancing method of the present invention as described herein. Additionally, performance of existing traffic on the path α - β will be improved as part of the load balancing method of the
10 present invention.

Consider now new traffic entering node B destined for node C. The natural path is via link β , which provides a direct path between nodes B and C; however, this link is heavily loaded with traffic from node A
15 destined for node C. Therefore, it may be better to send new incoming traffic via the path comprising links e, a, b, c (path e-a-b-c).

To account for traffic variations, load balancing is performed at periodic time intervals on a per class basis. For discussion, it will be assumed that traffic
20 flows are classified into three classes of service, including Class 1 (highest priority), Class 2 (medium priority), and Class 3 (lowest priority); however, it will be recognized that the teaching of the present
25 invention may be applied to any number of classes of service. Additionally, although for the sake of simplicity, the load balancing technique of the present invention will be illustrated as applied to Class 1, it

will be recognized that it may be applied to any class of service.

FIG. 2 is a block diagram of an existing QoS implementation within a single node 200. In general, as
5 illustrated in FIG. 2, packets entering the node 200 are classified by a packet classifier 222 into one of the three classes. Packets of each class are shaped by a respective token bucket policer 224 depending on various token bucket ("TB") parameters (e.g., peak data rate and
10 maximum burst size). Packets may be dropped if the traffic rate is excessive. Once shaped, packets are put into respective class based queues ("CBQs") 226 and drained from the CBQs 226 depending on the CBQ parameter specification of how many packets of each class can be serviced by a scheduler 228 at each service rotation.
15 The CBQ and TB parameters together comprise the quality of service ("QoS") parameters.

A bandwidth broker ("BB") 230 located at a higher layer (e.g., layer 3) is responsible for bandwidth
20 allocation to each class for a particular link. The BB parameters per class include allocated bandwidth, used bandwidth, and available bandwidth. The allocated bandwidth for a class refers to the total bandwidth available for that class on a link. Of this, only a
25 portion of the bandwidth is used by the various flows. The difference between allocated bandwidth and used bandwidth is the available bandwidth for that class on

that link. The sum of the allocated bandwidths for all classes equals the useful link capacity.

Assuming for the sake of example that the node 200 is node A of the RPR 100 (FIG. 1), in a top portion 200a of the node 200, the QoS parameters are applied to traffic on link a between nodes A and E and in a bottom portion 200b of the node 200, the QoS parameters are applied to traffic on link α between nodes A and B. Likewise for the BB parameters, even though the BB 230 is shown as a single entity for simplicity.

FIG. 3 is a block diagram of a node 300 embodying features of the present invention for implementing a load balancing technique. The node 300 is identical in all respects to the node 200 (FIG. 2), except that it includes a QoS/BB monitor 302. For purposes of clarity, only interaction of the monitor 302 with Class 1 is illustrated in FIG. 3; however, it should be recognized that, in practice, the monitor 302 is common to all rings in the RPR and monitors the traffic performance for each class on each ring at the node 300. With this knowledge, it can change the QoS parameters of each class on each link one at a time, for a particular traffic class. Likewise for the BB parameters.

Traffic performance (e.g., delay, jitter, dropped packets) is sensed by the QoS/BB monitor 302 at periodic time intervals by an in- or out-of-band signaling mechanism, such as that described in commonly-assigned U.S. Patent Application Serial No. _____ (Atty.

Docket No. 1285-0078US), filed _____, entitled
SIGNALING FOR CONGESTION CONTROL, LOAD BALANCING, AND
FAIRNESS IN A RESILIENT PACKET RING, which is hereby
incorporated by reference in its entirety. Changes to
the QoS and BB parameters are made after load imbalance
is detected, as will be described in greater detail
below.

The load balancing technique will now be described
with reference to FIGs. 1,3 and 4. Load balancing is
initiated either after a failure event or at periodic
time intervals T_k . The following variables, as defined
below, will be used to illustrate the load balancing
technique of the present invention. In connection with
a variable, subscript 1 denotes to Class 1, subscript 2
denotes Class 2, and subscript 3 denotes Class 3.
Additionally, subscript X denotes the clockwise ring
(comprising links α -e in FIG. 1) and subscript Y denotes
the counter-clockwise ring (comprising links a-e in FIG.
1). Notation is defined below only for Class 1.

B_{1X}, B_{1Y}	# bytes in CBQ drained off each rotation:
T_{p1X}, T_{p1Y}	TB peak rate
T_{b1X}, T_{b1Y}	TB burst size (in bytes)
T_k	time interval at which traffic measurement are made to detect load imbalance
T_{all1X}, T_{all1Y}	allocated bandwidth for links α and e, respectively
T_{av1X}, T_{av1Y}	available bandwidth for links α and e, respectively
T_{ul1X}, T_{ul1Y}	used bandwidth for links α and e, respectively

5 d_{1X}, d_{1Y} delays of the marked packets for Class 1
 obtained at specific time intervals T_k or when
 there is a failure event
 dr_{1X}, dr_{1Y} # dropped packets for Class 1 during the time
 interval T_k
 j_{1X}, j_{1Y} jitter of packets (to determine these
 variables, two successive packets for the
 class need to be marked)

10 It should be noted that available bandwidth and used
bandwidth are obtained from the BB entity 230 at higher
layers made available through an RPR control channel.
Example of constraints to which traffic in class 1 must
conform are defined as follows:

15 $dr_{1(X \text{ or } Y)} \leq DR_1$
 $d_{1(X \text{ or } Y)} + j_{1(X \text{ or } Y)} \leq D_{max1}$

 where DR_1 is the maximum acceptable number of dropped
packets for Class 1 and D_{max1} is the maximum acceptable
delay for Class 1. Similar constraints are defined for
20 other classes.

 FIG. 4 is a flowchart of the operation of a load
balancing method for RPR in accordance with teachings of
the present invention. Execution begins in step 400 at
a time T . In step 400, load imbalance detection is
25 performed and a determination is made whether a load
imbalance has been detected. Load imbalance detection at
each node is skewed with respect to other nodes at time
intervals of m , where $m < T_k$. Detecting load imbalance
may be accomplished in any number of ways, including, for
30 example, measurements of delays and/or dropped packets.
One method of detecting load imbalance is to measure and
compare delays for test packets sent from node A to a

node midway, for example, node C, on each of the two rings. Otherwise, some other node destination may be picked to test for load imbalance. If the delay and/or number of dropped packets is significantly higher on one ring, e.g., ring X, than the other ring, e.g., ring Y, for a given pair of nodes, then it could trigger the load balancing operation, as described below. Alternatively the constraints for each of the rings X and Y could be evaluated to determine a trigger for the load balancing operation. The threshold for triggering the load balancing operation should be sufficiently low so that the QoS/BB monitor can be sensitive to dynamic traffic patterns, but sufficiently high to prevent too frequent load balancing from occurring.

If load imbalance is not detected in step 400, step 400 is repeated again at the next time interval T_k (step 401). If load imbalance is detected in step 400, in step 402, responsive to a decision to implement load balancing, the QoS/BB monitor 302 signals the BB entity 230 to change the BB parameters so that new flows are diverted to ring Y rather than being sent via ring X. Thus, the BB parameters must be changed so that more bandwidth is available for new incoming Class 1 flows on ring Y and less bandwidth is available for new incoming traffic for Class 1 on ring X. Therefore, QoS/BB monitor 302 signals the BB monitor 230 to decrease T_{allX} and increase T_{allY} . This will effectively cause a Call Admission Control (CAC) entity (not shown), either co-

located with or separate from the BB entity 230, to admit more traffic flows on ring Y rather than on ring X, due to the increase in available bandwidth for Class 1 traffic on ring Y and the concurrent decrease in available bandwidth for Class 1 traffic on ring X.

Next, the QoS parameters for Class 1 on ring X and for Class 1 on ring Y must be changed so as to improve traffic performance on ring X while maximizing traffic performance on ring Y. This is accomplished as described below with respect to steps 404-414. In step 404, T_{p1X} is set to T_{u1X} . There is an interaction between the BB parameters and the QoS parameters. By signaling changes to the BB entity 230 to decrease T_{a11X} (step 402), fewer flows will be admitted on ring X; thus, T_{u1X} will decrease over a period of time, or not increase very much. In step 406, the TB parameters for Classes 2 and 3 on ring X (T_{p2X} , T_{b2X} , T_{p3X} , and T_{b3X}) are reduced. The reduction of these parameters relative to Class 1 traffic helps improve the performance of existing Class 1 traffic relative to other classes on ring X. If load imbalance is not subsequently detected, the Class 2 and Class 3 TB parameters can be increased again.

In step 408, T_{p1Y} is set to T_{u1Y} . Again, there is an interaction between the BB parameters and the QoS parameters. By signaling changes to the BB entity 230 to increase T_{a11Y} (step 402), more flows will be admitted on ring Y, and thus T_{u1Y} will increase. This helps improve performance of existing Class 1 traffic relative

to other classes of traffic on ring Y. In optional step 410, the TB parameters for Classes 2 and 3 on ring Y (T_{p2Y} , T_{b2Y} , T_{p3Y} , and T_{b3Y}) may be reduced to reduce the maximum Class 2 and Class 3 traffic rate/burst size serviced on ring Y. The actual amount of reduction or increase depends on the operator and its impact is described in greater detail below.

In step 412, B_{1X} is increased, while B_{2X} and B_{3X} are kept fixed. This preferentially permits more Class 1 traffic bytes to be serviced on ring X at each service rotation of the scheduler 228, relative to that for Classes 2 and 3, and thus helps improve the performance of Class 1 traffic. In step 418, B_{1Y} is increased, while B_{2Y} and B_{3Y} are kept fixed. The actual amount of increase may be performed in powers of two, for example, or using some other algorithm. Again, if load imbalance is not subsequently detected, the Class 1 CBQ parameter can be reduced on each of the two rings X and Y.

Changing the QoS parameters affects the performance of traffic already on the ring. Changing the BB parameters affects call admission for new traffic flows. Note that since these parameters can be changed a periodic time intervals or when there is a failure event, there is an opportunity to correct an incorrect selection of parameters at the next time interval T_k using new measurement data. Therefore, the actual algorithm used to increase or decrease the QoS/BM paymasters affects only the time it takes to accomplish load balancing, but

does not alter the objectives and results of load balancing.

Other options may also be implemented on the QoS/BB monitor. Various load imbalance detection algorithms can be selected. The selection of the measurement time interval is also left to the operator. The shorter the timer interval T_k , the more measurements that are available and thus, a better judgment of load imbalance can be made at the expense of computation time. This described mechanism for load balancing can easily be extended to RPRs comprising more than two rings or wavelengths.

Based upon the foregoing Detailed Description, it should be readily apparent that the present invention advantageously provides an innovative and efficient solution for correcting load imbalance in an RPR. In particular, because the QoS/BB monitor is common to all rings in an RPR, load balancing can be implemented easily, with minimal computation. The load balancing algorithm (FIG. 4) works by exercising control of traffic ingress via the BB parameters, and of traffic already on the RPR via the QoS parameters. The load balancing technique described herein allows correction of load imbalance on a class-by-class basis. Incorrect guesses for the QoS parameters and/or BB parameters can be corrected at the next time interval when new measurements are available. Finally, coordination between the QoS/BB monitors of various nodes can be implemented for other

load balancing, congestion control, and fairness algorithms.

It is believed that the operation and construction of the present invention will be apparent from the foregoing Detailed Description. While the exemplary
5 embodiments of the invention shown and described have been characterized as being preferred, it should be readily understood that various changes and modifications could be made therein without departing from the scope of
10 the present invention as set forth in the following claims.

WHAT IS CLAIMED IS:

1 1. A method of implementing load balancing in a
2 resilient packet ring ("RPR") network comprising a
3 plurality of nodes and first and second rings each
4 comprising a plurality of links for carrying information
5 between the nodes in a clockwise direction and a
6 counterclockwise direction, respectively, wherein
7 adjacent ones of the nodes are connected by two of the
8 links, the method comprising the steps of, for one of the
9 nodes:
10 determining whether a load imbalance exists at the
11 node in connection with a first class of service; and
12 responsive to a determination that a load imbalance
13 exists:
14 changing Bandwidth Broker ("BB") parameters at
15 the node for the first class of service to cause new
16 flows to be diverted from a more heavily loaded one of
17 the rings to a less heavily loaded one of the rings; and
18 changing Quality of Service ("QoS") parameters
19 at the node for the first class of service to improve
20 traffic performance on the more heavily loaded one of the
21 rings, while maximizing traffic performance on the less
22 heavily loaded one of the rings.

1 2. The method of claim 1 wherein the step of
2 determining is performed at periodic time intervals.

1 3. The method of claim 1 wherein the step of
2 determining is performed using a technique selected from
3 the group consisting of measuring and comparing delays
4 experienced by a test packet sent from the node to a
5 second node via the first and second rings, respectively,
6 and comparing a number of dropped packets on the first
7 and second rings with a preselected maximum value.

1 4. The method of claim 1 further comprising the
2 step of signaling to a QoS/BB monitor that a load
3 imbalance has been detected responsive to a determination
4 that a load imbalance exists.

1 5. The method of claim 1 wherein the step of
2 changing the BB parameters comprises the steps of:
3 decreasing an allocated bandwidth for the first
4 class of service on the more heavily loaded ring; and
5 increasing an allocated bandwidth for the first
6 class of service on the less heavily loaded ring.

1 6. The method of claim 1 wherein the step of
2 changing the QoS parameters comprises the steps of:
3 setting the peak traffic rate to the used bandwidth
4 for the first class of service on the more heavily loaded
5 ring;
6 reducing token bucket ("TB") parameters for all
7 other classes of service on the more heavily loaded ring;
8 setting the peak traffic rate to the used bandwidth
9 for the first class of service on the less heavily loaded
10 ring; and
11 increasing the number of bytes in a class based
12 queue ("CBQ") for the first class of service drained off
13 in each scheduler rotation for each of the rings.

1 7. The method of claim 6 wherein the step of
2 changing the QoS parameters further comprises the step of
3 reducing TB parameters for all other classes of service
4 on the less heavily loaded ring.

1 8. Apparatus for implementing load balancing in a
2 resilient packet ring ("RPR") network comprising a
3 plurality of nodes and first and second rings each
4 comprising a plurality of links for carrying information
5 between the nodes in a clockwise direction and a
6 counterclockwise direction, respectively, wherein
7 adjacent ones of the nodes are connected by two of the
8 links, the apparatus comprising, at one of the nodes:
9 means for detecting at the node a load imbalance in
10 connection with a first class of service;
11 means responsive to detection at the node of a load
12 imbalance for changing Bandwidth Broker ("BB") parameters
13 at the node for the first class of service to cause new
14 flows to be diverted from a more heavily loaded one of
15 the rings to a less heavily loaded one of the rings; and
16 means responsive to detection at the node of a load
17 imbalance for changing Quality of Service ("QoS")
18 parameters at the node for the first class of service to
19 improve traffic performance on the more heavily loaded
20 one of the rings, while maximizing traffic performance on
21 the less heavily loaded one of the rings.

1 9. The apparatus of claim 8 wherein the means for
2 detecting performs the detecting at periodic time
3 intervals.

1 10. The apparatus of claim 8 wherein the means for
2 detecting comprises means for measuring and comparing
3 delays experienced by a test packet sent from the node to
4 a second node via the first and second rings.

1 11. The apparatus of claim 8 wherein the means for
2 detecting comprises means for comparing a number of
3 dropped packets on the first and second rings with a
4 preselected maximum value.

1 12. The apparatus of claim 8 further comprising
2 means for signaling to a QoS/BB monitor that a load
3 imbalance has been detected.

1 13. The apparatus of claim 8 wherein the means for
2 changing the BB parameters comprises:

3 means for decreasing an allocated bandwidth for the
4 first class of service on the more heavily loaded ring;
5 and

6 means for increasing an allocated bandwidth for the
7 first class of service on the less heavily loaded ring.

1 14. The apparatus of claim 8 wherein the means for
2 changing the QoS parameters comprises:

3 means for setting the peak traffic rate to the used
4 bandwidth for the first class of service on the more
5 heavily loaded ring;

6 means for reducing token bucket ("TB") parameters
7 for all other classes of service on the more heavily
8 loaded ring;

9 means for setting the peak traffic rate to the used
10 bandwidth for the first class of service on the less
11 heavily loaded ring; and

12 means for increasing the number of bytes in a class
13 based queue ("CBQ") for the first class of service
14 drained off in each scheduler rotation for each of the
15 rings.

1 15. The apparatus of claim 14 wherein the means for
2 changing the QoS parameters further comprises means for
3 reducing TB parameters for all other classes of service
4 on the less heavily loaded ring.

1 16. Apparatus for implementing load balancing in a
2 resilient packet ring ("RPR") network comprising a
3 plurality of nodes and first and second rings each
4 comprising a plurality of links for carrying information
5 between the nodes in a clockwise direction and a
6 counterclockwise direction, respectively, wherein
7 adjacent ones of the nodes are connected by two of the
8 links, the apparatus comprising, at one of the nodes:
9 a Quality of Service/Bandwidth Broker ("QoS/BB")
10 monitor responsive to detection at the node of a load
11 imbalance in connection with a first class of service for
12 signaling to a BB to change BB parameters at the node for
13 the first class of service to cause new flows to be
14 diverted from a more heavily loaded one of the rings to
15 a less heavily loaded one of the rings and for changing
16 QoS parameters at the node for the first class of service
17 to improve traffic performance on the more heavily loaded
18 one of the rings, while maximizing traffic performance on
19 the less heavily loaded one of the rings.

1 17. The apparatus of claim 16 wherein detection of
2 a load imbalance is accomplished by measuring and
3 comparing delays experienced by a test packet sent from
4 the node to a second node via the first and second rings.

1 18. The apparatus of claim 16 wherein detection of
2 a load imbalance is accomplished by comparing a number of
3 dropped packets on the first and second rings with a
4 preselected maximum value.

1 19. The apparatus of claim 16 wherein the QoS/BB
2 monitor is apprised of a load imbalance via an in-band
3 signaling mechanism.

1 20. The apparatus of claim 16 wherein the QoS/BB
2 monitor is apprised of a load imbalance via an out-of-
3 band signaling mechanism.

1 21. The apparatus of claim 16 wherein the QoS/BB
2 monitor changes the BB parameters by:
3 decreasing an allocated bandwidth for the first
4 class of service on the more heavily loaded ring; and
5 increasing an allocated bandwidth for the first
6 class of service on the less heavily loaded ring.

1 22. The apparatus of claim 16 wherein the QoS/BB
2 monitor changes the QoS parameters by:
3 setting the peak traffic rate to the used bandwidth
4 for the first class of service on the more heavily loaded
5 ring;
6 reducing token bucket ("TB") parameters for all
7 other classes of service on the more heavily loaded ring;
8 setting the peak traffic rate to the used bandwidth
9 for the first class of service on the less heavily loaded
10 ring; and
11 increasing the number of bytes in a class based
12 queue ("CBQ") for the first class of service drained off
13 in each scheduler rotation for each of the rings.

1 23. The apparatus of claim 22 wherein the QoS/BB
2 monitor changes the QoS parameters by reducing TB
3 parameters for all other classes of service on the less
4 heavily loaded ring.

ABSTRACT OF THE DISCLOSURE

A technique for load balancing in an resilient packet ring ("RPR") is disclosed. In one embodiment, the present technique comprises implementing on every node an QoS/BB monitor, which is common to all rings in the RPR and has knowledge of traffic performance for each class on each ring of the RPR, which information is obtained through periodic measurements or in response to failure events. This allows the monitor to vary the QoS parameters on a node, for a particular traffic class, to achieve load balancing. Likewise, the QoS/BB monitor can signal to the BB entity at higher layers to vary the BB parameters on a node for a particular class to achieve load balancing.




FIG. 1




FIG. 2
(Prior Art)




FIG. 3

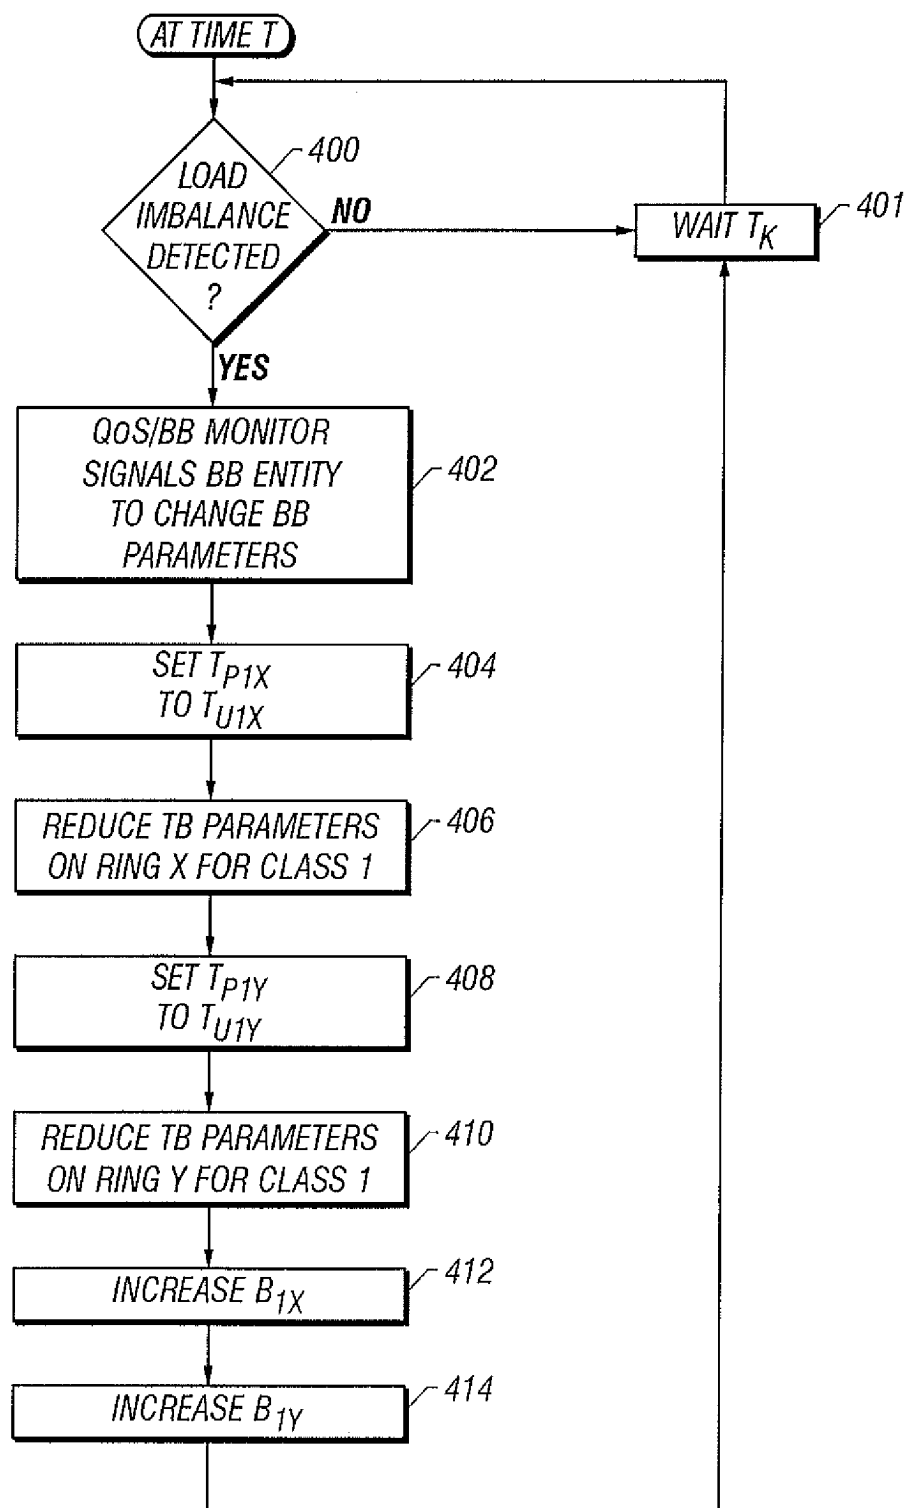


FIG. 4